P15914

APPLICATION FOR UNITED STATES LETTERS PATENT

INVENTORS: Edward A. BURTON and Kumar ANSHUMALI

TITLE: ADDITION OF METAL LAYERS WITH SIGNAL REALLOCATION TO A MICROPROCESSOR FOR INCREASED FREQUENCY AND LOWER POWER

ATTORNEYS & ADDRESS: FLESHNER & KIM, LLP
P. O. Box 221200
Chantilly, VA 20153-1200

DOCKET NO.: (Intel) 2112-021

FILING DATE: June 27, 2003

ASSIGNEE: INTEL CORPORATION

ADDITION OF METAL LAYERS WITH SIGNAL REALLOCATION TO A MICROPROCESSOR FOR INCREASED FREQUENCY AND LOWER POWER

BACKGROUND OF THE INVENTION

1. Field of the Invention

[1] The present invention relates to semiconductor design. In particular, the present invention relates to adding metal layers to a microprocessor to increase operating frequency and lower power consumption.

2. Background of the Related Art

[2] Related art microprocessor designs have an interconnect stack in which signal wires in alternate metal layers are orthogonal to each other. A perspective block diagram is shown in Figure 1, which illustrates the related art design. In interconnect stack is illustrated with layers 110, 120, 130 and 140. Metal layers 120 and 130 are illustrated with signal wires arranged in an orthogonal pattern as known in the related art, while metal layers 110 and 140 have been approximated by a plane.

[3] Figure 2 illustrates an axial view of a related art metal layer interconnect stack. Ground planes 202 and 204, in this example, are above and below layer 200, respectively. Layer 200 includes a plurality of signal wires including power supplies 210 at right and left ends, ground/return 212 and signal wire sets 220 and 230. Each signal wire set includes a plurality of signal wires 221-226 and 231-234, respectively. The first set 220 of six signal wires 221-226 is at a minimum pitch. The second set 230 of four signal wires 231-234 is at a

20% larger pitch. Those skilled in the art will appreciate that the illustrated arrangement is solely for purposes of illustration.

[4] The related art topology creates problems in regard to the parasitic capacitance (i.e., the wire capacitance) on the signal wires. The wire capacitance decreases the potential operating frequency and increases the power consumed due to the capacitive load. Since the operating frequency of semiconductor devices is constantly increasing, reducing the wire capacitance of the signal wires allows for increased frequency and reduced power consumption, particularly on high frequency signal wires that are limited in the related art architecture.

BRIEF DESCRIPTION OF THE DRAWINGS

[5] The invention will be described in detail with reference to the following drawings in which like reference numerals refer to like elements wherein:

Figure 1 illustrates a perspective view of a related art interconnect stack containing a plurality of signal wires and layers;

Figure 2 illustrates an axial view of a related art metal layer containing a plurality of signal wires;

Figure 3 illustrates a system level block diagram of a computer system according to an embodiment of the present invention;

Figure 4 illustrates a perspective view of an interconnect stack containing a plurality of signal wires and layers according to an embodiment of the present invention;

Figures 5A, 5B, 5C and 5D illustrate axial views of metal layers containing a plurality of signal wires according to embodiments of the present invention;

Figure 6A illustrates an example of identifying a signal wire as a critical signal wire;

Figures 6B and 6C illustrates process parameters with regard to signal wires; and

Figures 7A, 7B and 7C are flowcharts for adding metal layer containing a plurality of signal wires according to embodiments of the present invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

[6] In the following detailed description of preferred embodiments, reference is made to the accompanying drawings that show, by way of illustration, specific embodiments in which the invention may be practiced. In the drawings, like numerals describe substantially similar components throughout the several views. These embodiments are described in sufficient detail to enable those skilled in the art to practice the invention. Other embodiments may be utilized, and structural, logical, and intellectual changes may be made without departing from the scope of the present invention. Moreover, it is to be understood that various embodiments of the invention, although different, are not necessarily mutually exclusive. For example, a particular feature, structure, or characteristic described in one embodiment may be included within other embodiments. The following detailed description is not to be taken in a limiting sense, and the scope of the present invention is defined only by appended claims, along with the full scope of equivalence to which such claims are entitled.

[7] Embodiments of the present invention relate to signal wire placement and design improvements for use in a wide variety of applications including computer systems. Figure 3 shows an exemplary illustration of a computer system. The computer system can include a microprocessor 2, which can include an interconnect stack with additional metal layers and signals reallocated to improve the maximum frequency and/or power requirements. Microprocessor 2 can include many sub-blocks such as an arithmetic logic unit (ALU) 4 and an on-die cache 6. The microprocessor 2 may also communicate to other levels of cache, such as off-die cache 8. Higher memory hierarchy levels such as system memory 10 are accessed via host bus 12 and a chip set 14. In addition, other off-die functional units such as a graphics accelerator 16 and a network interface controller 18, to name just a few, may communicate with the microprocessor 2 via appropriate busses or ports. Examples of an interconnect stack with additional metal layers and signals reallocated are presented in the following description.

[8] Figure 4 illustrates a perspective view of an arrangement of an interconnect stack according to an embodiment of the present invention. The illustration can be understood in relation to the illustration of the related art arrangement of Figure 1. As shown in Figure 4, the interconnect stack includes layers 410, 420, 430, and 440, and two additional layers 420' and 430' that have been added to metal layers 420 and 430. The additional layers 420' and 430' preferably are orientated in the same direction as layers 420 and 430, respectively. Additionally, the signal wires of the related art layer 120 in Figure 1 are distributed over layers 420 and 420'. Likewise, the signal wires of the related art layer 130 in Figure 1 are distributed over layers 430 and 430'. As is apparent from Figure4, the

signal wire separation is increased and consequently the wire capacitance for each signal wire is reduced since the signal wires are distributed over two layers. The reduced wire capacitance results in a digital gate driving that wire to have a lower load. In addition the RC delay of the wire decreases because of the lowered capacitance in the signal wire.

[9] Figure 5A illustrates an axial view of an arrangement of an interconnect stack according to an embodiment of the present invention. This illustration can be understood in relation to the illustration of the related art arrangement of Figure 2. Ground planes 502 and 504, in this example, are above and below metal layer 500 and additional metal layer 500', respectively. Layers 500 and 500' include a plurality of signal wires including power supply wires 510 at right and left ends, ground/return wire 512 and signal wire sets 520 and 530. Each signal wire set 520, 530 includes a plurality of signal wires 521-526 and 531-534, respectively. The power supply wires 510 and ground/return wire 512 are located in both layers 500 and 500'. Although the illustrated embodiment has the power supply and ground wires in both layers, they do not have to be in both layers. However, those skilled in the art will appreciate that if a power or ground wire is present to satisfy a stringent shielding requirement (e.g., for noise reduction, delay matching, and the like), the power and/or ground lines next to the shielding-critical signal must be replicated. Information regarding the shielding requirements for a given signal can be accessed, for example, from a design database for the device as is known in the art.

[10] The first set 520 of six signal wires 521-526 is distributed among layers 500 and 500'. Particularly, the signal wires 521-526 are alternated between the layers with signal wires 521, 523, and 525 located in layer 500' and signal wires 522, 524, and 526 located in

layer 500. When compared to the arrangement illustrated in Figure 2, it is apparent that alternating the signal wires between the base 500 and prime 500' layers increases the separation of the signal wires 521-526.

[11] Likewise, the second set of signal wires 530 is distributed among layers 500 and 500'. Specifically, the signal wires 531-534 are alternated between the layers with signal wires 531 and 533 located in layer 500' and signal wires 532 and 534 located in layer 500. Those skilled in the art will appreciate that this arrangement is shown solely for purposes of illustration and not for limitation of the invention. The added metal layer (i.e., the prime layer 500') preferably has similar process parameters to the original metal layer (e.g., 500) in the device. However, in contrast to conventional orthogonal arrangements, the orientation of the signal wires of the prime layer is substantially parallel to the original layer. In addition, the number of signal wires on either layer can be reduced after a prime layer is added and the signals are reallocated.

[12] Accordingly, an embodiment of the present inventions includes a semiconductor device that has a first layer and a second layer. The first and second layers are adjacent and each have a plurality of signal wires, as illustrated in Figure 5A, for example. The signal wires in the first and second layers are substantially parallel with each other. As illustrated in Figure 5A, adjacent signal wires are located in a substantially alternating pattern in the first and second layers. For example, every other signal wire (521, 523, 525, 531, 533) is located in the second layer and remaining signal wires (522, 524, 526, 532, 534) are located in the first layer. Therefore, the substantially alternating pattern has every other signal wire residing in the same layer in this illustration. However, not all signal wires have to follow

this alternating pattern. For example, signal wires can be dedicated to specific layers or to both layers for various design purposes such as interconnections, shielding and the like. Even if all signal wires are not located in an alternating pattern, the reduced wire capacitance due to the additional metal layer and relocation of some of the signal wires still can improve the performance of the device.

[13] Further, the power supply 510 and ground 512 wires can be located in both the first layer and second layer and in the same respective position in each layer, as illustrated. Therefore, the alternating pattern stops at the power and ground wires. For example, referring to Figure 5A, the alternating pattern stops at either the power supply wires 510 or the ground/return wires 512, because the power and ground wires are located in both layers.

[14] The foregoing description illustrates the specific embodiment of Figure 5A. Another arrangement of the signal wires according to an embodiment of the present invention is illustrated in Figure 5B. For convenience, the signal wires, layers, power supply and ground wires are labeled the same as in Figure 5A. Referring to Figure 5B, an axial view of an arrangement of an interconnect stack according to another embodiment of the present invention is presented. Once again, this illustration may be best understood in relation to the illustration of the related art arrangement of Figure 2. Ground planes 502 and 504, in this example, are above and below metal layer 500 and additional metal layer 500', respectively. Layers 500 and 500' include a plurality of signal wires including power supplies 510 at right and left ends, ground/return 512 and signal wire sets 520 and 530. Each signal wire set 520, 530 includes a plurality of signal wires 521-526 and 531-534, respectively. In

this embodiment, the power supply wires 510 and ground/return wire 512 are located only in layer 500. Additionally, signal wire set 530 has the signal wires 531-534 replicated in each layer, as opposed to the alternating pattern of set 520. Those skilled in the art will appreciate that certain signals in a design (e.g., word lines) can be connected to devices that have high capacitance. In this case, reducing the wire capacitances of these signals may not have the greatest impact on the RC component of the signal path. Therefore, reducing the resistance of the signal wire can be used to reduce the power losses and RC component in the signal wires thus improving the performance of these signal paths.

[15] An arrangement of signal wires according to another embodiment of the present invention is illustrated in Figure 5C. For convenience, the signal wires, layers and power and ground wires are labeled the same as in Figure 5A. In this embodiment the signal wire set 530 is not replicated or arranged in an alternating pattern. Instead, wire set 540 is arranged in an orthogonal pattern with respect to the signal wires of the first layer. The first set 520 of six signal wires 521-526 is distributed among layers 500 and 500'. Specifically, the signal wires of set 520 are alternated between the layers with signal wires 521, 523, and 525 located in layer 500' and signal wires 522, 524, and 526 located in layer 500 as in the previous illustrations. However, in this configuration, areas of the device that do not have frequency limiting signal paths (e.g., set 530) can have signal wires arranged orthogonal to the signal wires of the first layer.

[16] Additionally, those skilled in the art will appreciate that the design benefits are not limited to adding only one additional metal layer. For example, as illustrated in Figure 5D, two metal layers have been added 500' and 500", when compared to Figure 2. The

signal wires 521-526 and 531-534 have been distributed over the layers to reduce the wire capacitance as in the previous examples. Figure 5D also illustrates that the distribution of the signal wires does not have to result in the same pattern, as illustrated in signal wire sets 520 and 530. Accordingly, multiple additional layers can be added. Each additional layer also has signal wires substantially parallel to signal wires in the first and second layers. Once again, adjacent signal wires can be located in an alternating pattern or replicated in the first, second and additional layers.

[17] In addition, more than one layer of signal wires in a device can benefit from adding an additional layer. Accordingly, an embodiment of the present invention includes a semiconductor further comprising a third and fourth layer having a plurality of signal wires. For example, layers 420 and 420' are the first and second layers and layers 430 and 430' are the third and fourth layers, as illustrated in Figure 4. The signal wires in the third and fourth layers are substantially parallel with each other and adjacent signal wires are distributed in the third and fourth layers, for example, using the alternating pattern illustrated. However, those skilled in the art will appreciate that any of the foregoing embodiments or combinations thereof can also be used to distribute the signal wires in the third and fourth layers. Providing the additional metal layer allows for the signal wires to be distributed between the third and fourth layers, thus reducing the wire capacitance and/or resistance of the signal wires.

[18] In general, there are two issues that drive the decision on how many additional layers are inserted. First, the preponderance of speed-limiting paths on a layer or layers.

Second, the preponderance of power contributing paths on a layer or layers. These two issues are discussed in more detail in the following sections.

[19] Regarding the first issue, an N-metal layer system can have speed limiting paths evenly distributed on all layers. In this case, the maximum benefit is obtained from adding N extra (i.e., prime) layers to the design, one for each pre-existing layer. However, in a typical microprocessor design, for example, the lower N/2 metal layers (or more) are utilized for local interconnection (e.g., within a functional block such as an adder, incrementer, and the like). The upper N/2 metal layers can carry global routes such as system clocks, long-distance signals, power grid, and the like. Accordingly, the upper N/2 layers are more likely to contain the frequency limiting speed paths. Depending on the speed path distribution with respect to the metal layers, the prime layer can be inserted only on the original layers that had enough speed paths to justify the additional cost of adding the prime layer. Those skilled in the art will appreciate that the justification of adding the prime layers is empirical and based on the additional processing cost, market demand, frequency improvement, power reduction, and the like.

[20] Regarding the second issue, the power contributing paths are not necessarily the same as the frequency limiting speed paths, since the power aspect is dependent on total capacitance and activity level on the path. However, once the power contributing paths are identified reducing the wire capacitance of the signal wires or the effective resistance by replicating signal wires can reduce the power consumption, as discussed in the foregoing sections.

[21] Given the above considerations, the layers in a design that have the most impact on the speed limiting paths and high power consumption should be identified. The speed-limiting paths (e.g., due to capacitive loading and/or RC delay) can be identified by a timing analysis of the whole chip with the interconnect capacitances and resistances of signal paths on various layers modeled accurately, as is known in the art. The power consumption can be determined by monitoring nodal activity while running a high power test pattern (e.g., in a microprocessor design), as is known in the art. Layers that have a significant contribution to the overall power can then be identified. Based on the identified speed limiting and high power consuming layers, additional layers can be inserted adjacent to these layers to reduce the impact on the semiconductor device's performance in terms of increased operating frequency and/or reduced power consumption.

[22] Figure 6A illustrates a plurality of signal wires on a layer that embodiments of the present invention can be applied to. For example, signal wires 600 can be analyzed and the critical signal wire 601 on the layer is used as a seed for a method according to an embodiment of the present invention. The method will be described in the following sections.

[23] Additionally, the layers (i.e., the prime layers) that are added preferably have the same process parameters. Figure 6B illustrates a signal wire and related process parameters in the original layer. Wires in a process technology (e.g., 0.13 μm) have a thickness "t", which can be different for different layers and a minimum width W_{min}, which is the minimum width at which a wire can be drawn. Once again, different layers can have different minimum widths. However, an inserted layer (i.e., prime layer) preferably supports

the same thickness and minimum width as the original layer. Further, wires in a layer can have different permittivity for sidewall capacitance (ε_{side}) versus up capacitance (ε_{up}) versus down capacitance (ε_{down}). Accordingly, the inserted layer (i.e., prime layer) preferably has similar dielectric permittivities to its neighbors as the original. This can be difficult to accomplish if ε_{up} is not equal to ε_{down}. However, in this case ε_{up} can be made the same on the upper (e.g., prime layer 500') and ε_{down} can be made the same on the lower layer (e.g., original layer 500). An example of this arrangement is illustrated in Figure 6C. Although the process parameters are not the same with respect to the permittivity at the border of the layers, the increased spacing can still reduce the wire capacitance.

[24] Figure 7A is a flow chart illustrating a method of adding metal layers in a semiconductor device with alternating signal wire reallocation according to an embodiment of the present invention. In block 710, a second metal layer is added to a first metal layer. The first and second layers are adjacent. A plurality of signal wires is distributed between the first layer and the second layer, in block 712. The plurality of signal wires is preferably substantially parallel with each other. As discussed above, adding the second layer (i.e., prime layer) allows for increased separation of the signal wires and/or replication of the signal wires to reduce the wire capacitance and/or resistance and consequently the RC delay and power consumption.

[25] Figure 7B illustrates additional steps regarding distributing the plurality of signal wires. The critical signal wire from a plurality of signal wires in the first layer is identified, in block 720. For example, the critical signal wire can be a signal wire that is frequency limited due to the wire capacitance in a conventional layer design, as discussed in

the foregoing description. Accordingly, the frequency limiting speed paths are identified and the critical signal wire is used as a seed to start the signal wire reallocation process.

[26] The signal reallocation process can start by assigning the signal wires adjacent to the critical signal wire to the second layer (or alternate layer), in block 730. The first layer is checked to see if there are remaining signal wires to reallocate, in decision block 740. The process continues assigning the remaining signal wires in an alternating pattern to the second layer, in block 750 until all signal wires are assigned. For example, starting at the critical signal wire, the signal wires on the left and right are assigned to the second layer. Then, the process can continue in both directions. Thus, the next signal wire encountered in both directions would stay in the first layer, and the signal wire after that would be moved to the second layer and so on until all signal wires have been processed.

[27] Figure 5A illustrates the result of applying the forgoing method to a signal wire arrangement as illustrated in Figure 2. The configuration in Figure 2 showed all signal wires in a first layer 200. In Figure 5A, the signal wires have been distributed between the first layer 500 and the second layer 500' in an alternating pattern. Additionally, note that the power supply 210 and ground 212 wires in Fig 2 are replicated in each layer 500 and 500' in Figure 5A. Therefore, the foregoing method can further include inserting power 510 and ground 512 wires into the second layer, as illustrated in optional block 715. Additionally, the power and ground wires in the second layer can have the same respective position (e.g., horizontal position) as the power and ground wires in the first layer, as illustrated in Figure 5A. Accordingly, a method according to an embodiment of the present invention can further comprise stopping the alternating pattern at power supply and ground wires and then

resuming the alternating pattern at a next signal wire after the power supply or the ground. This aspect is illustrated in Figure 7C.

[28] Referring to Figure 7C, an embodiment of block 750 is illustrated. In decision block 752, the next wire is checked to see if it is a power or ground wire. If a power or ground wire is encountered, then that wire is skipped (i.e., the alternating pattern stops). However, if the next wire is a signal wire, then the pattern continues and assigns that signal wire to a different layer than the previous signal wire (i.e., the alternating pattern continues). If the previous wire were a power supply or ground, then the pattern would resume at the next signal wire following the power supply or ground wire. Regardless of whether a power supply, ground or signal wire is encountered, the process returns and loops to block 740 to check for remaining signal wires to assign.

[29] Once again, the foregoing description in regard to Figures 7B and 7C illustrates specific examples of embodiments of the present invention and has been provided solely for illustration and not limitation. Alternative methods can be appreciated from the foregoing illustrations of Figures 5A-5D. For example, distributing the signal wires can include replicating a set of signal wires in both the first and second layers, as illustrated in Figure 5B. Further, additional signal wires in the second layer can be arranged substantially orthogonal to the signal wires in the first layer, as illustrated in Figure 5C.

[30] Those skilled in the art will appreciate that embodiments of the present invention are not limited to only two layers. Accordingly, a method according to an embodiment of the present invention can further comprise providing at least one additional layer, such as illustrated in Figure 5D. The signal wires in the at least one additional layer are

substantially parallel signal wires in the first and second layers. Further, adjacent signal wires can be located in an alternating pattern in the first, second and at least one additional layer.

[31] Additionally, embodiments of the present invention can be used to improve the performance of existing microprocessor designs. For example, in Figure 7B block 720, identifying the critical signal wire, can further include analyzing an existing microprocessor design to identify frequency limiting signal wires, using techniques known in the art, such as those discussed in the foregoing description. Then, one of the frequency limiting signal wires is identified as the critical signal wire. In the case when there is more than one frequency limiting signal wire, the most critical (e.g., the signal wire that is most limited) is used as the critical signal wire. Once the critical signal wire in the existing (first) layer is identified, the signals are distributed over the second layer and first layer to reduce the wire capacitance and consequently improve the potential frequency and power consumption of the existing design. Alternate methods of assigning an alternating pattern to signal wires can also be used. For example, even a random selection the signal wires and/or patterns will result in some improvement because of the increased and reduced wire capacitance. Additionally, critical signal wires can be identified as the signal wires in a layer that have high power consumption, as discussed in the foregoing description.

[32] As noted with respect to Figure 3, embodiments of the present invention relate to signal wire placement and design improvements for use in a wide variety of applications including computer systems. Accordingly an embodiment of the present invention includes a computer system comprising a microprocessor 2 and an off-die component (e.g., 8, 10, 12, 18, 16) in communication with the microprocessor. The

microprocessor comprises a first layer having a plurality of signal wires and a second layer having a plurality of signal wires. The signal wires in the first and second layers are substantially parallel with each other and adjacent signal wires are located in an alternating pattern in the first and second layers. Further, the various embodiments described in the foregoing description can be applied to the microprocessor and other components in the system to increase the potential operating frequency and reduce the power consumption of these devices. Additionally, those skilled in the art will appreciate that the term microprocessor used herein refers broadly to processing devices including microprocessors, Digital Signal Processors (DSPs), microcontrollers, and the like

[33] The foregoing descriptions of embodiments of the present invention discuss the benefit of reducing parasitic capacitance, with respect to frequency. Additionally, the reduction of the wire capacitance reduces the dynamic load. For example, a reduction of the wire capacitance in the interconnect stack of ten percent, will result in an approximately ten percent reduction in the dynamic power consumed in the interconnect stack.

[34] Further, consider a microprocessor with fifty percent of its power loss coming from wire capacitance. By inserting as many additional metal layers as there are in an existing microprocessor design (e.g., 2 to 4, or from 4, to 8), the interconnect contribution to the power loss can be reduced by half. Still further, the drivers on the inputs can be reduced in size because of reduced wire capacitance, leading to approximately a twenty-five percent or more reduction in power.

[35] The foregoing embodiments and advantages are merely exemplary and are not to be construed as limiting the present invention. The present teaching can be readily

applied to other types of apparatuses. The description of the present invention is intended to be illustrative, and not to limit the scope of the claims. Many alternatives, modifications, and variations will be apparent to those skilled in the art.